

Efficient Parallelization for AMR MHD Multiphysics Calculations; Implementation in AstroBEAR

Jonathan J. Carroll-Nellenback^a, Brandon Shroyer^a, Adam Frank^a, Chen Ding^b

^a*Department of Physics and Astronomy, University of Rochester, Rochester, NY 14620*

^b*Department of Computer Science, University of Rochester, Rochester, NY 14620*

Abstract

Current AMR simulations require algorithms that are highly parallelized and manage memory efficiently. As compute engines grow larger, AMR simulations will require algorithms that achieve new levels of efficient parallelization and memory management. We have attempted to employ new techniques to achieve both of these goals. Patch or grid based AMR often employs ghost cells to decouple the hyperbolic advances of each grid on a given refinement level. This decoupling allows each grid to be advanced independently. In AstroBEAR we utilize this independence by threading the grid advances on each level with preference going to the finer level grids. This allows for global load balancing instead of level by level load balancing and allows for greater parallelization across both physical space and AMR level. Threading of level advances can also improve performance by interleaving communication with computation, especially in deep simulations with many levels of refinement. To improve memory management we have employed a distributed tree algorithm that requires processors to only store and communicate local sections of the AMR tree structure with neighboring processors.

1. Introduction

The development of Adaptive Mesh Refinement [3, 4] methods were meant to provide high resolution simulations for much lower computational cost than fixed grid methods would allow. The use of highly parallel systems and the algorithms that go with them were also meant to allow higher resolution simulations to be run faster (relative to wall clock time). The parallelization of AMR algorithms, which should combine the cost/time savings of both methods is not straight forward however and there have been many different approaches [6, 9, 8, 1, 2]. While parallelization of a uniform mesh demands little communication between processors, AMR methods can demand considerable communication to maintain data consistency across the unstructured mesh as well as shuffling new grids from one processor to another to balance workload.

In this paper we report the development and implementation of new algorithms for the efficient parallelization of AMR designed to scale to very large simulations. The new

Email address: johannjc@pas.rochester.edu (Jonathan J. Carroll-Nellenback)

algorithms are part of the AstroBEAR package for simulation of astrophysical fluid multi-physics problems [5]. The new algorithmic structure described in this paper constitutes the development of version 2.0 of the AstroBEAR code.

AstroBEAR like many other grid based AMR codes utilizes a nested tree structure to organize each individual refinement region. However, as we will describe, unlike many other AMR codes, AstroBEAR 2.0 uses a distributed tree in which no processor has access to the entire tree but rather each processor is only aware of the AMR structure it needs to manage in order to carry out its computations and perform the necessary communications. While currently, this additional memory is small compared to the resources typically available to a CPU, future clusters will likely have much less memory per processor similar to what is already seen in GPU's. Additionally each processor only sends and receives the portions of the tree necessary to carry out its communication.


AstroBEAR 2.0 also uses extended ghost cells to decouple advances on various levels of refinement. As we show below this allows for each level's advance to be computed independently on separate threads. Such inter-level threading allows for total load balancing across all refinement levels instead of balancing each level independently. Independent load balancing becomes especially important for deep simulations (simulations with low filling fractions but many levels of AMR) as opposed to shallow simulations (high filling fractions and only a few levels of AMR). Processors with coarse grids can advance their grids simultaneously while processors with finer grids advance theirs. Without such a capability, each level would need to have enough cells to be able to be distributed across all of the processors. Variations in the filling fractions from level to level can make the number of cells on each level very different. If there are enough cells on the level with the fewest to be adequately distributed, there will likely be far too many cells on the highest level to allow the computation to be completed in a reasonable wall clock time. This often restricts the number of levels of AMR that can be practically used. With inter-level threading this restriction is lifted. Inter-level threading also allows processors to remain busy while waiting for messages from other processors.

In what follows we provide descriptions of the new code and its structure as well as providing tests which demonstrate its effective scaling. In section 2 we review patch based AMR. In section 3 we will discuss the distributed tree algorithm, in section 4 we will discuss the inter-level threading of the advance, in section 5 we will discuss the load balancing algorithm, and in section 6 we will present our scaling results and we will conclude in section 7.

2. AMR Algorithm

Here we give a brief overview of patch based AMR introducing our terminology along the way. The fundamental unit of the AMR algorithm is a patch or grid. Each grid contains a regular array of cells in which the fluid variables are stored. Grids with a common resolution or cell width Δx_l belong to the same level l and on all but the coarsest level are always nested within a coarser "parent" grid of level $l-1$ and resolution $\Delta x_{l-1} = R \times \Delta x_l$ where R is the refinement ratio. The collection of grids comprises the AMR mesh, an example of which is shown in figure 1. In addition to the computations required to advance the fluid variables, each grid needs to exchange data with its parent grid (on level $l-1$) as well as any child grids (on level $l+1$). Grids also need to exchange data with physically adjacent neighboring grids (on level l). In order to exchange data,

Figure 1: Example AMR mesh showing nested and adjacent grids as well as corresponding AMR tree showing parent-child and neighbor relationships.




the physical connections between parents, children, and neighboring grids are stored in the AMR tree as connections between nodes. Each grid has a corresponding node in the AMR tree. Thus there is a one to one correspondence between nodes and grids. The grids hold the actual fluid dynamical data while the nodes hold the information about each grid's position and its connections to parents, children and neighbors. Figure 1 shows one example of an AMR mesh made of grids and the corresponding AMR tree made of nodes. Note that what matters in terms of connections between nodes is the physical proximity of their respective grids. While siblings share a common parent, they will not necessarily be neighbors, and neighbors are not always siblings but may be 1st cousins, 2nd cousins, etc...

Additionally since the mesh is adaptive there will be successive iterations of grids on each level as the simulation progresses. Thus the fluid variables need to be transferred from the previous iteration of grids to the current iteration as shown in figure 2. Thus nodes can have “neighbors” within a 4 dimensional space-time. Nodes that are temporally adjacent (belonging to either the previous or next iteration) and spatially coincident are classified as preceding or succeeding overlaps respectively instead of temporal neighbors, reserving the term neighbor to refer to nodes of the same iteration that are spatially adjacent and temporally coincident. Nodes on level l therefore have a parent connection to a node on level $l - 1$, child connections to nodes on level $l + 1$, neighbor connections to nodes on level l of the same iteration, and overlap connections to nodes on level l of the previous or successive iteration in time.

2.1. Regridding

We use the term “iteration” to refer to successive generations of the grid distribution for each level. Thus at some point in the simulation a distribution of level l grids will

Figure 2: Example AMR mesh showing two iterations of level 1 grids and the overlapping regions. Nodes associated with previous/successive iterations of overlapping grids will have preceding/succeeding overlap connections in the AMR tree.



cover some fraction of the computational space. When, as the simulation proceeds, the grid generation subroutine is called again the next iteration of level l grids will be laid down over the computational domain.

While in principle level l could be regridded each level l time step, it is convenient for restricting data to parent grids, to wait until after R level l steps, or equivalently after 1 parent level ($l - 1$) step so that child and parent grids have advanced to the same time. Therefore for each level 0 time step, there is 1 iteration of level 1 grids, R iterations of level 2 grids, R^2 iterations of level 3 grids, and so on. Since new child grids are created after each parent step, each parent will have multiple iterations of children. This additional branching of the AMR tree in time allows for temporally adjacent (overlap) grids to be classified as temporal siblings or 1st, 2nd, 3rd temporal cousins etc... Additionally parents will have connections to multiple iterations of children (although only the two most current need to be kept).

Note that there are different ways of indexing the iteration of grids on a given level. One can chose a global indexing that begins with the initiation of the simulation, i.e the iterations can be indexed by counting the number of successive iterations of grids on that level from the beginning of the simulation. Another method is to index iterations relative only to the level above (IE which iteration of children from a given parent does a grid correspond to.) We will use the latter indexing scheme, since it is relevant to the way overlap connections between grids of different iterations are formed.

Figure 3 shows this indexing system (where we have assumed $R = 2$ for simplicity). Here the circles represent iterations of the entire level (i.e. the entire collection of nodes for a given level at the specified iteration). Since figure 3 shows the evolution of the AMR tree *in time*, it's useful to make a connection with the the way the tree appears at

any given moment *in space* in terms of spatial connections between nodes (figure 1). To understand the relation between figure 1 and figure 3 imagine taking the tree in figure 1 and rotating it into the plane. Neighbor relationships would now disappear and all of the nodes of a given level would visually merge. Each visible circle in figure 3 represents a collection of nodes of the same level and iteration as shown in figure 1. Each level ($l > 0$) has multiple iterations that stretch both forward and backward in time. Note that the level 0 grids are static and there are no successive iterations. Level 1 iterations go from 1 to ∞ since the level 0 grids continue to create successive iterations of children. Level 2 iterations and higher are indexed from 1 to $R = 2$.

It is important to stress here, that each level's iteration in figure 3 represents a collection of nodes on that level. Additionally each level's iteration represents multiple time steps on that level. The nodes of the level 0 iteration actually take many time steps, while all other level iterations take $R = 2$ time steps. Preceding overlap connections (those going backward in time) are only needed before the 1st of R time steps, and the succeeding overlap connections (those going forward in time) are only needed after the last of R time steps. After the 1st time step and before the last, ghost overlap data is shared between neighboring grids of the same iteration instead of preceding/succeeding overlaps.

While there will be many iterations of grids on each level, only the two most current iterations are stored in the tree. After a grid finishes its advance it becomes "old". When this occurs the previous generation of old grid iterations are discarded. The current and old grids are shown schematically in figure 3. The current level l iteration will always contain children of the current level $l - 1$ iteration. For the old level l iteration there are two cases. First the old level l iteration may be the old iteration of children of the current $l - 1$ iteration as is the case for the old level 3 grids in figure 3. Secondly, the old level l iteration may be the last iteration of children of the old level $l - 1$ iteration as is the case for the old level 2 grids in figure 3. That is, they are either old children of the current parent level grids, or last children of the old parent level grids.

2.2. Grid Life Cycle

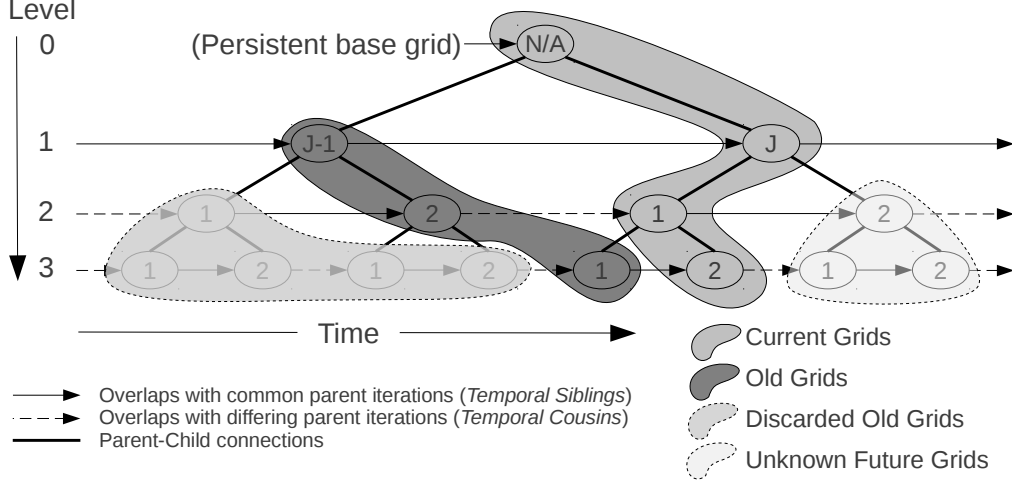
The execution of the various stages in the AMR algorithm is shown in the upper panel of Figure 5 (for $R = 2$). There are 5 steps in the algorithm: overlaps (O); Prolongation (P); Advances (A); Synchronization (S) and Restriction (R). A pseudo-code description of their order of operations with L being the highest level of refinement is written below. Note each root step is initialized by calling $\text{AMR}(0)$.

```

RECURSIVE SUBROUTINE AMR( $l$ )
  Overlap data from preceding iterations of grids on level  $l$ 
  DO  $i=1, R$ 
    Advance grids on level  $l$ 
    IF ( $l < L$ ) THEN
      Prolongate initial (pre-advanced) data from level  $l$  to level  $l + 1$ 
      CALL AMR( $l + 1$ )
      Restrict data from level  $l + 1$ 
    END IF
    Synchronize fluxes between neighboring level  $l$  grids and update ghost cells
  END DO

```

Figure 3: Time dependence of AMR tree showing overlaps between old grids and new grids. Each node is labeled by the iteration of its parent's children.



END RECURSIVE SUBROUTINE

The subroutine above describes a process in which grids are created, populated with data, advanced and brought in line with higher and lower refinement representations of the data. To be more specific, after each iteration of level l grids are created their cells (including ghost regions) are initialized with a combination of prolonged data from the parent grid as well as data from the preceding set of level l grids that physically overlap with the computational space that the new grids now describe. Ghost regions are needed to update the fluid variables within the grid. The new level l grids then determine which cells to refine and lay down a first set of level $l + 1$ child grids to cover those cells. The parent level l grids take one step of Δt_l . Meanwhile their child grids advance R steps of $\Delta t_{l+1} = \Delta t_l / R$. Next the level l grids merge restricted data from their children with their own updated data since both levels have now reached the same time t . These level l grids then synchronize fluxes with any adjacent neighboring grids (also of level l) and update any ghost cells. The level l grids are now ready to repeat the process (creating children; advance a time step; restricted child data; synchronize fluxes with neighbors). When the level l grids have completed R steps relative to the level above ($l - 1$), their own data and fluxes are restricted and applied to their parent ($l - 1$) grids.

After completing its advances, the data from each level l grid is stored until it can be copied onto the next iteration of level l grids. After the data is copied onto *succeeding* overlaps, the old level l grids are destroyed. Throughout a level l grid's lifetime it must therefore share data with its parent grid ($l - 1$), its preceding overlaps (l), multiple iterations of its own child grids ($l + 1$), its neighbor grids (l), and succeeding overlaps (l). These connections for a given grid are held by its node and the web of connections between grids/nodes form the AMR tree. While a grid/node may have many successive iterations of children, only connections that belong to the two most current iterations of the child level are kept.

2.3. Isolated Grids

Refined grids need to take multiple (R) time steps. For the first step, prolonged ghost cells from the coarser grid can be used to advance the internal cells. However further time steps cannot be taken without updating the ghost cells in time as well. If a grid is surrounded by neighboring grids, it can update its ghost cells in time by copying the advanced data from its neighbors' internal cells. However, if a grid is isolated (ie. has no neighbors) or even if it is partially isolated, it must advance those ghost cells in another way. There are two solutions to this problem shown in figure 4. One is to use the time derivative for the fluid variables calculated on the coarse grid to update the ghost cells where needed. This solution requires parent grids to advance before their children and can lead to large errors if shocks are not completely resolved since spatial discontinuities lead to delta functions for time derivatives which, in turn, produce large errors upon discretization.

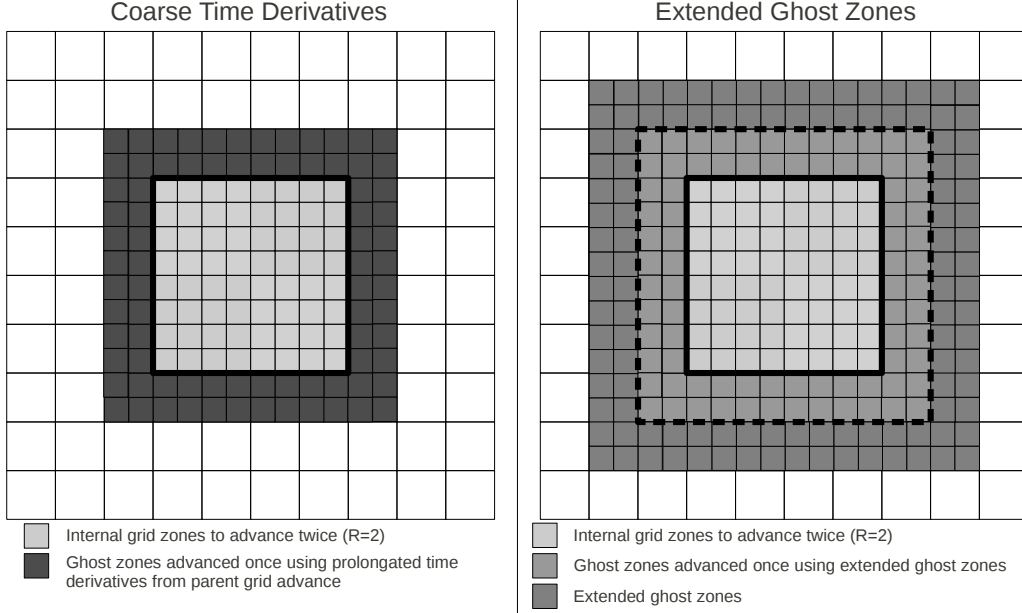
The alternative method is to use extended ghost cells that allow grids to successively update smaller and smaller regions so that there is always enough ghost cells to complete the final step. In general if n_{ghost} is the number of ghost cells needed to take a time step and R is the refinement ratio, then the extended ghost region needs to initially be $R \times n_{\text{ghost}}$ cells wide. While only isolated grids need to update extended ghost cells, it can be difficult to implement efficient advance schemes for partially isolated grids, and each grid is often assumed to be isolated. While this can result in a large overhead for small grids, (especially when the refinement ratio R , or n_{ghost} is large) the ability to allow coarse grids to be advanced independently of fine grids can be exploited to increase the speed of the code. We have taken this approach in AstroBEAR to allow for more efficient distribution of grids as well as the creation of multiple advance threads described in section 4.

3. Distributed Tree Algorithm

Many current AMR codes store the entire AMR tree on each processor. This, however, can become a problem for simulations run on many processors. To demonstrate this let us first assume that each AMR grid requires m bytes per node to store its meta data (ie 6 bytes to store its physical bounds for a 3D simulation and 1 byte to store the processor containing the grid). We also assume that each grid requires on average d bytes for the actual data. If there are, on average, n grids on each of p processors, then the memory per processor would be $nd + nmp$. The second term nmp represents the meta-data for the *entire* AMR tree.

The memory requirement for just the nodes in the AMR tree without storing any connections becomes comparable to the local actual data when $p = d/m$. If we assume a 3D isothermal hydro run where each cell contains $\rho, p_x, p_y, \& p_z$ with a typical average grid size of $8 \times 8 \times 8$ then $p = \frac{8 \times 8 \times 8 \times 4}{(6+1)} \approx 293$. While this additional memory requirement is negligible for problems run using typical CPU's on 100's of processors, it can become considerable when the number of processors $n > 10^3$. Since it is expected that both efficient memory use and management will be required for ever larger HPC (high performance computing) clusters down the road, AstroBEAR 2.0 is designed to use a *distributed* tree algorithm. In this scheme each processor is only aware of the section of the tree containing nodes that connect to its own grids' nodes. Additionally, new nodes

Figure 4: Two solutions for isolated grids. The one on the right requires additional calculations to update the ghost cells once, but allows for each level of the AMR hierarchy to advance independently.



are communicated to other processors on a proscriptive “might need to know basis”. Since maintaining these local trees as the mesh adapts is not trivial, we describe the process below.

3.1. Maintaining the AMR Tree

Because of the nested nature of grids, neighbor and overlap relationships between nodes can always be inherited from parent relationships. First consider the neighbors of the n^{th} iteration of a node’s children. The nested nature of the grids restricts each of the child’s neighbors to either be a sibling of that child (having the same parent node and be of the same iteration), or to be a member of a neighbor’s n^{th} iteration of children. Thus the neighbors of a level l node’s children (on level $l + 1$) will always be a child of that level l node’s neighbors.

For overlaps (temporal neighbors) the situation is a bit more complicated. If we consider figure 3 we can identify two different types of overlap connections. First there are overlap connections between iterations with a common parent which we will refer to as ‘temporal siblings’. For example the current and old iterations shown on level 3 are different iterations of children of the same level 2 iteration. There are also overlap connections between children that come from different iterations of the parent grids. These can be thought of as ‘temporal cousins’. For example the current and old child iterations shown on level 2 come from different iterations of the parent nodes one level above them. In principle temporal cousins can be 1st cousins, 2nd cousins, and so on, but the distinction does not matter for the way connectivity is treated in the algorithm.

Child iteration	Node's child's neighbors	Node's child's preceding overlaps	Node's child's succeeding overlaps
1^{st}	Node's neighbors' 1^{st} children	Node's preceding overlaps' 2nd children	Node's neighbor's 2^{nd} children
2^{nd}	Node's neighbors' 2^{nd} children	Node's neighbors' 1^{st} children	Node's succeeding overlaps' 1^{st} children

Overlaps between temporal siblings on level l will be between successive child iterations of the current level $l - 1$ iteration. Note that grids in the current $l - 1$ iteration do not physically overlap. Since grids are always nested, the overlaps between the grid's current (n^{th}) iteration of children must belong to the grid's old ($n^{th} - 1$) iteration of children (for completeness these are what are referred to as preceding overlaps). However, while grids do not physically overlap, their ghost cells may. Thus overlapping data for a level l grid's children may also come from a neighbor level l grid's old children. Thus the (preceding) overlaps for a node's children may be an old child of the node's neighbors. Likewise, every node's old child's succeeding overlap may be a node's neighbor's current child.

For temporal cousins, the current iteration of level $l - 1$ grids/nodes will be the first iteration of children of the current level l grids/nodes. And the old iteration of level $l - 1$ grids will be the last iteration of children of the old level l grids. Because of the nested nature of the grids, we can say that for the first iteration of children, every node's child's preceding overlap must be a node's preceding overlap's child. The same follows for the succeeding overlaps of a node's last iteration of children. That is, every node's child's succeeding overlap must be a node's succeeding overlap's child. These relationships are summarized for $R = 2$ in table 3.1 and a generalized table is given in the appendix.

3.2. Maintaining Distributed Sub-Trees

For parallel applications, the grids are distributed across the processors. In addition to data for the local grids, each processor needs to know where to send and receive data for the parents, neighbors, overlaps, and children of those local grids. This is information contained within the nodes. In order for a processor to know where to send data, each one must maintain a local sub-tree containing its own "local" nodes (corresponding to local grids) as well as all remote nodes (living on other processors) directly connected to the local nodes. It is also possible, though not desirable, that an individual processor have data from disjoint regions of the simulation. In that case each processor would have multiple disjoint sections of the AMR tree, but these disjoint sections would collectively be considered the processor's sub-tree.

Each time new grids on level $l + 1$ are created (by local parent grids on level l), each processor determines how the new child grids should be distributed (i.e. which processor should get the new grids). This distribution is carried out in the manner described in section 5 below. Connections between the new level $l + 1$ nodes and the rest of the tree must then be formed. Because of the inheritability of the neighbor/overlap connections, even if a child grid is distributed to another processor, the connections between that child node and its neighbors/overlaps/parent are first made on the processor that created the grid (ie the processor containing the grid's parent). If a processor's local grid has a remote parent, then that processor will always receive information about that local grid's neighbors/overlaps from the processor containing the remote parent. This is true

of neighbor and preceding overlap connections of new grids as well as succeeding overlap connections of old grids.

Before processors containing parents of level $l + 1$ grids (both new and old) can send connection information to remote children (if they exist), these processors must first share information about the creation of children with each other. Neighbor connections between new nodes on level $l + 1$ require each processor to cycle through its local level l nodes and identify those with remote neighbors living on other processors. Once remote neighbors have been identified the information about new children from the local nodes is sent to the processor(s) containing the remote neighbors. Not all children need to be sent to every remote neighbor. Only those that are close enough to potentially be adjacent to the remote neighbor's children are necessary. The information must flow in both directions meaning a individual processor also needs to *receive* information about potential new children from all other processors containing remote neighbors.

Overlaps are again a bit more complicated because of the different situations for temporal siblings and temporal cousins as described above. For temporal cousins (ie if the new level $l + 1$ child grids are the 1st iteration from level l parents), each processor must cycle through local current level l nodes with preceding overlaps that live on remote processors. The processor must then send new children of local nodes to the processor(s) containing the remote preceding overlaps. Similarly, each processor must cycle through the local **old** level l nodes with remote succeeding overlaps sending children of those old local nodes to the processor(s) containing remote succeeding overlaps. Again, not all children need to be sent to the processors containing remote overlaps. Only information about those nodes close enough to potentially overlap the remote node's children must be sent. The information must flow in both directions and processors also need to receive information about potential children from those same processors with remote overlaps.

For overlaps between temporal siblings, no communication is required. This is because previous level $l + 1$ nodes that could overlap a new level $l + 1$ child node would either be old children of the same parent or be old children of the parent's neighbors. The parent would already be aware of its old children, and the old children of the parent's neighbors would have been previously communicated when establishing the neighbor relationships of the parent's old children. This process is summarized in the appendix. table .1

4. Threaded Multilevel Advance

Many if not all current AMR codes tend to perform grid updates across all levels in a prescribed order that traverses the levels of the AMR hierarchy in a sequential manner. Thus the code begins at the base grid (level 0), moves down to the highest refinement level and then cycles up and down across levels based on time step and synchronization requirements (for a simulation with 3 levels the sequence would be: 0, 1, 2, 2, 1, 2, 2, 0...) In the top panel of figure 5 the basic operations of (P)rolongating, (O)verlapping, (A)dvancing, (S)ynchronizing, and (R)estricting are shown for each level along with the single (serial) control thread.

Good parallel performance requires each level update to be independently balanced across all processors (or at least levels with a significant fraction of the workload). Load balancing each level, however, requires the levels to contain enough grids to be effectively distributed among the processors. Such a requirement demands each level to be fairly "large" in the sense of having many grids or allowing each level's spatial coverage be

artificially fragmented into small pieces. The former situation leads to broad simulations (large base grid leaving resources for only a few levels of AMR), while the later situation leads to inefficient simulations due to the fair amount of overhead required for ghost cell calculations.

The problem becomes worse when n_{ghost} is large or when using extended ghost cells with refinement ratios $R > 2$. Consider an isolated level l grid of size 4×4 inside of a parent grid on level l . Let's assume that the coarsening ratio $R = 2$ and that the grid must therefore take two steps each of which requires $n_{ghost} = 4$ ghost cells. On the first step it must update a region that is 12×12 and then on the second step a region that is 4×4 for a total of $12 \times 12 + 4 \times 4 = 160$ cell updates. If that 4×4 grid were split into $4 \times 2 \times 2$ pieces to be distributed on 4 processors, then instead of 160 cell updates there would now be $10 \times 10 + 2 \times 2 = 104$ per processor or 416 total cell updates and the parallelization efficiency for that grid would be $160/416 = 38\%$ at best. Even if one uses prolonged time derivatives instead of extended ghost cells, there is still a fair amount of overhead for ghost cell computations for small grids and the efficiencies would still be $\approx 60 - 70\%$ depending on the particular method used.

In the bottom panel of figure 5 we show a schematic of the AstroBEAR 2.0 AMR algorithm. In this figure basic operations of (P)rolongating, (O)verlapping, (A)dvancing, (S)ynchronizing, and (R)estricting are shown again however this time the level advances are independent and exist on separate threads of computation. There is an overarching control thread which handles all of the communications and computations required for prolongation, overlapping, synchronizing, and restricting as well as the finest level advances. Each coarser level advance has its own thread and can be carried forward independently with preference being given to the threads that must finish first (which is always the finer level threads). In addition to relaxing the requirement of balancing every level, the existence of multiple threads allows processors to remain busy when the control thread becomes held up because it needs information from another processor. For example, while waiting for ghost cell data for level 3 it can work on advancing levels 2, 1, or 0.

The simplest implementation of threading would allow for the operating system kernel to manage the various threads with priority given to the control thread followed by the highest level advance and so on all on the same processor. This would be considered preemptive prioritized threading within a process. Unfortunately the Pthreads library implementation under Linux does not allow for non-privileged users to give priorities to threads nor to limit a set of threads to use only a single processor. GNU portable threads do require the threads to remain on a single processor, but do not allow for preemptive thread scheduling. Threads must schedule themselves cooperatively by yielding control to the scheduler or to each other. While GNU portable threads do allow for priorities to be assigned to threads, it is not terribly useful as threads can cooperatively yield control to the appropriate thread without the need for a scheduler. A more complicated implementation, but one that does not require external libraries, is to manually switch (at the application level) between grid advances and the control thread. This is somewhat difficult to implement as the grid advances must be interruptable and any intermediate variables within the advance stack have to be manually cached by the application. This caching effectively limits the frequency with which the advance thread can be interrupted. In AstroBEAR 2.0 we have implemented GNU portable threads, manual thread switching, as well as the serial version of the AMR algorithm. Performance results are

Variable List	
g_l^p	Current workload assigned to processor p on level l
w_l^p	Amount of current workload g_l^p completed
c_l^p	Newly created child workload on level l by parent grids on level $l - 1$ on processor p
$\overline{g_l}$	Average of g_l^p across processors
$\overline{w_l}$	Average of w_l^p across processors
G_l	Sum of g_l^p across processors. Also equals sum of c_l^p across processors.
d_l^p	Desired workload assignments for processor p and level l
ϵ_l^p	Imbalance in g_l^p ($g_l^p - \overline{g_l}$)
ω_l^p	Imbalance in w_l^p ($w_l^p - \overline{w_l}$)
η_l^p	Predicted remaining workload for entire AMR step for processor p on level l
s_l	Current number of remaining steps on level l within entire AMR step
$\delta_{p'}^p$	Workload assigned to processor p' of children created by grids on processor p

given in section 6.

5. Load Balancing

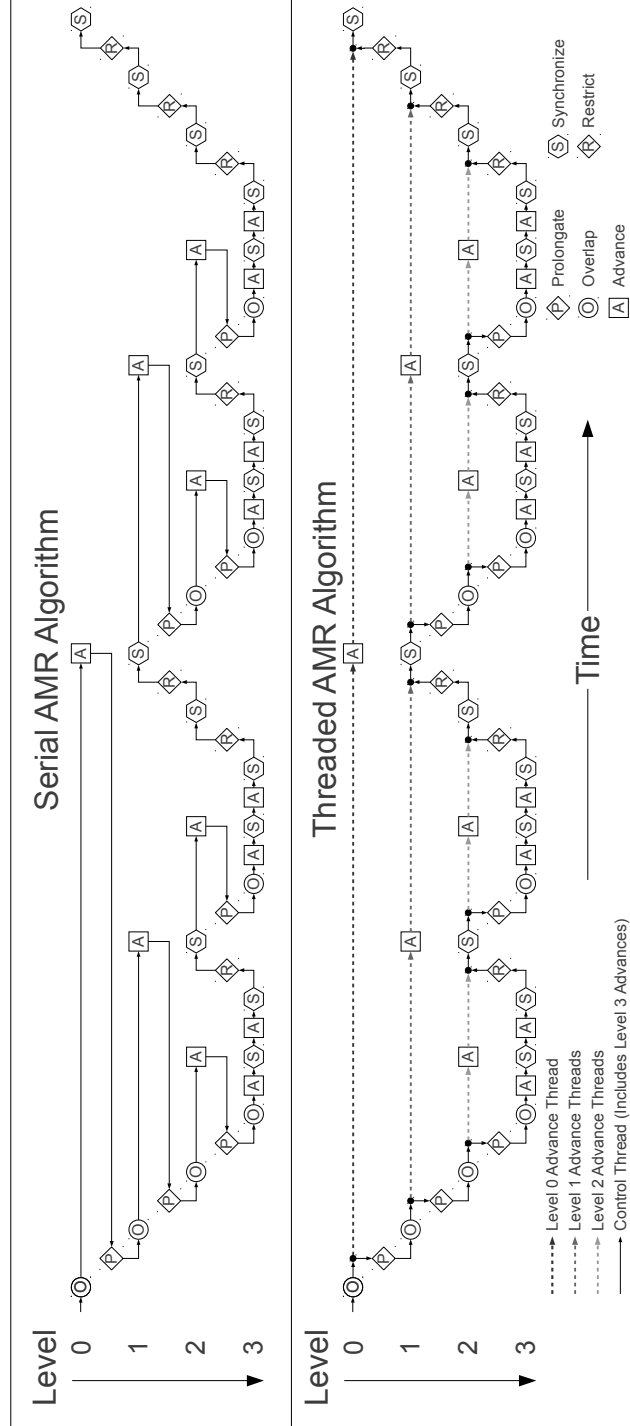
As was discussed above, threading level advances removes the need for balancing each level independently and instead allows for global load balancing. It also and perhaps more importantly allows for consideration of the progress of coarser advance threads when successively distributing the workload of finer grids. This “dynamic load balancing” allows adjustments to be made to finer level distributions to compensate for variations in progress made on ongoing coarser advances.

Here we give a brief example introducing terminology along the way. Also see table 5 for a complete list of variables used below. Consider a simulation run on three processors with two levels of refinement with a refinement ratio $R = 2$ and let us assume (arbitrarily) that there are $G_0 = 72$ cells on level 0 distributed among the three processors as $g_0^p = [32, 20, 20]$. Each processor may have more than one grid, but here we just count the total number of cells in all of the processor’s grids. After spawning the level 0 advance thread, each processor continues along the control thread until it eventually creates new child grids on level 1. At this point these new child grids must be distributed. Let’s assume that at this point no work has yet been accomplished on the level 0 advance thread ($w_0^p = [0, 0, 0]$). We also assume that the number of new child cells created by grids on each processor is $c_1^p = [32, 8, 8]$ for a total of ($G_1 = 48$) cells created on level 1 to be distributed across all three processors. If we were independently balancing the workload on level 1 the desired distribution d_1^p would be constant for each processor

$$d_1^p = \overline{g_1} \quad (1)$$

where $\overline{g_1} = \overline{c_1} = \frac{G_1}{N}$ is the average new child workload. For our example this would give a distribution of $d_1^p = [16, 16, 16]$. There is however a workload imbalance on level 0 of $\epsilon_0^p = g_0^p - \overline{g_0} = [+8, -4, -4]$ that we would like to compensate for. In what follows we show how we can successively rewrite d_1^p in new forms which allow us to better anticipate the best distribution to account for computation within and across levels.

Figure 5: Plot showing threads of AMR algorithm



If we completely compensate for the level 0 imbalance

$$d_1^p = \bar{g}_1 - \epsilon_0^p \quad (2)$$

then the desired workload distribution for level 1 would be $[8, 20, 20]$. If the level 1 workload is distributed that way, then while processor 1 waits for processors 2 & 3 to complete their 20 level 1 updates, it should have time to complete 12 of its level 0 updates so $w_0^p = [12, 0, 0]$. At this point the remaining workload on level 0 would be balanced at $g_0^p - w_0^p = [20, 20, 20]$. However we need to take an additional step on level 1. Thus while the remaining level 0 workload is balanced, we see an opportunity to redefine d_1^p in a way that accounts for this additional level 1 step. We would like to redistribute the grids on level 1 such that they are balanced at $[16, 16, 16]$ before taking the second step to avoid processor idling. This redistribution has an additional cost that can easily be avoided if the imbalance on level 0 is first weighted by the ratio of the number of level 0 steps s_0 to level 1 steps s_1 remaining. During the first distribution of level 1 grids there are 2 remaining level 1 steps $s_1 = 2$ and one remaining level 0 step $s_0 = 1$. Thus we redefine $d - 1^p$ by modifying equation 2 to

$$d_1^p = \bar{g}_1 - \frac{s_0 \epsilon_0^p}{s_1} \quad (3)$$

This would give us $[16, 16, 16] - [+4, -2, -2] = [12, 18, 18]$. Now while processor 1 waits for processors 2 & 3 to complete their 18 level 1 updates, it should have time to complete 6 of its level 0 updates effectively reducing the imbalance by a factor of 2. After completing the first level 1 advance, the second distribution of level 1 grids must somehow take into account the coarser workload accomplished. To do so we include a completed workload imbalance $\omega_l^p = w_l^p - \bar{w}_l = [4, -2, -2]$ into equation 3.

$$d_1^p = \bar{g}_1 - \frac{s_0 \epsilon_0^p - \omega_0^p}{s_1} \quad (4)$$

Note that now with $s_1 = 1$ we get the same desired distribution of $[12, 18, 18]$ for the second level 1 step as for the first.

We could have just used the remaining workload $h_0^p = g_0^p - w_0^p$ to calculate a remaining workload imbalance $\epsilon_0^p = h_0^p - \bar{h}_0$ and then used equation 3 but this would be the same as modifying equation 4 to be $d_1^p = \bar{g}_1 - \frac{s_0(\epsilon_0^p - \omega_0^p)}{s_1}$. Since in our example $\omega_0^p \neq 0$ only when $s_0 = 1$, it makes no difference, but when there are more than two levels of refinement this approach can lead to unnecessary shuffling of grids between processors even with static meshes. For a more detailed analysis see the appendix.

When there are multiple levels of refinement, the workloads on many coarser levels must be taken into account. The generalization of equation 4 is given by

$$d_l^p = \bar{g}_l - \frac{\sum_{l'=0}^{l-1} s_{l'} \epsilon_{l'}^p - \omega_{l'}^p}{s_l} \quad (5)$$

A more convenient form of equation 5 can be obtained if we replace ϵ_l^p and ω_l^p with their expansions and then group terms by whether they involve averages across processors

or local quantities on a particular processor. We then arrive at the final general form of the distribution equation

$$d_l^p = \bar{g}_l - \frac{\eta_{l-1}^p - \bar{\eta}_{l-1}}{s_l} \quad (6)$$

where $\eta_{l-1}^p = \sum_{l'=0}^{l-1} s_{l'} g_{l'}^p - w_{l'}^p$. Note that $\bar{\eta}_{l-1}$ is the predicted remaining work on all coarser levels and that if the actual distribution matches the desired distribution (ie $g_l^p = d_l^p$) then the total predicted remaining workload including the current level is

$$\eta_l^p = \eta_{l-1}^p + s_l g_l^p = \eta_{l-1}^p + s_l d_l^p = \eta_{l-1}^p + s_l \bar{g}_l - \eta_{l-1}^p + \bar{\eta}_{l-1} = s_l \bar{g}_l + \bar{\eta}_{l-1} = \bar{\eta}_l. \quad (7)$$

That is, each distribution attempts to equalize the predicted remaining work over the entire AMR advance. It is not always possible to distribute a discrete set of grids perfectly so that $g_l^p = d_l^p$, however small differences on coarser levels can be corrected for on finer level distributions. On the finest level the grids are artificially split to balance out all of the coarser level imbalances as described below.

Thus at each distribution of level l grids, processors must collect the new child workloads c_l^p to calculate $\bar{g}_l = \bar{c}_l$, and η_l^p to calculate the desired workload distributions d_l^p . Then the new child workloads c_l^p are partitioned over $d_l^{p'}$ as shown in figure 6 to give the desired workload $\delta_{p'}^p$, assigned to processor p' of children created by grids on processor

p . Note that $\sum_{p'=1}^P \delta_{p'}^p = c_l^p$ and $\sum_{p=1}^P \delta_{p'}^p = d_l^{p'}$. Each processor p can then determine from

which "parent" processors p' it should expect to receive new child grids from ($\delta_{p'}^{p'} > 0$) as well as which "child" processors p'' it should distribute new grids to ($\delta_{p''}^p > 0$). Here the term child processor does not refer to the addition of new processors (unlike the addition/creation of new child grids) but only those existing processors which will receive child grids from other processors (referred to here as "parents").


When distributing child grids among child processors, each processor p will try to assign to each child processor p' a collection of level l grids with a combined workload $\delta_{p'}^p$. However, arbitrarily partitioning a few discrete grids into exact sizes is often not possible, and the actual workload assigned $\delta_{p'}^{p,*}$ to each child processor may be different from the desired workload $\delta_{p'}^p$. As a result, the combined child workload from all of a

processor's parent processors $g_l^p = \sum_{p'=1}^P \delta_{p'}^{p,*}$ may be different from the desired workload d_l^p . This results in some variation in η_{l+1}^p between processors and a temporary predicted load imbalance. However, if this variation is small, it can be corrected on finer level distributions or on the next round of level l distributions without adversely effecting performance, provided that there remains enough computational work on coarser advance threads to buffer the imbalance.

5.1. Grid Assignments

Before discussing the details of the distribution of new refined grids, it is instructive to first consider the simpler problem of distributing newly refined cells as is the case for

Figure 6: Example partitioning $\delta_p^{p'}$ of new level l child workloads $c_l^{p'}$ onto desired workload shares d_l^p . Shades of gray correspond to processor rank (0-3).



cell-based AMR. After newly refined cells are created, a space filling curve is generally used to sort the new cells into a prescribed order based on their position along the curve [7]. In addition the processors are ordered based on the topology of the network. Then the ordered list of cells can be partitioned over the ordered list of processors in a manner similar to that used in constructing the δ_p^p in figure 6 except that now instead of globally distributing c_l^p over d_l^p to determine δ_p^p , we are locally distributing the individual workload of each child cell c_l^i over the child processor allocations δ_p^i to determine the processor allocations for each individual child cell $\mathcal{D}_{p'}^i$. Most cells will have only one nonzero allocation $\mathcal{D}_{p'}^i > 0$ and will be assigned to that processor p' and if a cell has more than one nonzero allocation then the processor with the largest allocation could be used. This “rounding” will only lead to load imbalances that are of order $\frac{N_{processors}}{N_{cells}}$. This ordering keeps communication between neighboring cells to a minimum. For the partitioning to work, however, each processor needs to know the global sort index of each of its newly refined cells. Sorting the new cells could be done using traditional sorting algorithms provided each processor was aware of all of the cells. A better alternative, is to maintain a strict grouping of cells by sort order on each processor as shown in figure 7. That is, a cell with a lower sort index will never be found on a higher numbered processor and vice versa. Because of the fractal nature of many space filling curves like the Hilbert curves shown in figure 7, the ordering of any two child cells of different parent cells will be the same as that of its parents. Therefore if the parent cells obey a strict grouping by sort order, the child cells will as well. Processors therefore only need to determine the local sort order of their child cells. They can then determine the global sort order if they know the new child cell counts of every other processor. An example of these cell counts can be seen in figure 7 as well as local and global index of each level 0 and level 1 cells. The global index of each cell can be constructed by adding the cell counts of lower processors to the local index. The fact that child cells will obey the same ordering of their parents can be used to locally sort child cells by first sorting the parent cells and then locally sorting the child cells among their siblings.

In cell based AMR each processor can usually partition its newly refined cells among “child processors” with very little load imbalance since each processor typically has 1000’s of cells to distribute over a few processors. Not being able to split a cell might result in imbalances of few tenths of a percent. This is not the case for patch based AMR. Processors will still have 1000’s of cells, but they will be grouped into a few grids of various sizes. Distributing a few grids of various sizes among a few child processors

with various workload assignments will in general lead to significant imbalances. Often what is done in grid-based AMR is a more global distribution of all new child grids among all processors using a knapsack algorithm. This type of distribution does not rely on a space-filling curves or any global ordering of new grids although it may take the communication costs into account. This however, requires global knowledge of the AMR tree which can result in memory issues for simulations run on many processors. In AstroBEAR, the AMR tree is distributed so a knapsack type algorithm is not feasible. Instead AstroBEAR uses an approach quite similar to that used for cell-based AMR. AstroBEAR avoids the issues with load imbalances not by artificially fragmenting grids into small pieces, or by locally shuffling the grids in a local knapsack algorithm, but instead through the use of level threading described above in which load imbalances can be corrected for by subsequent distributions.

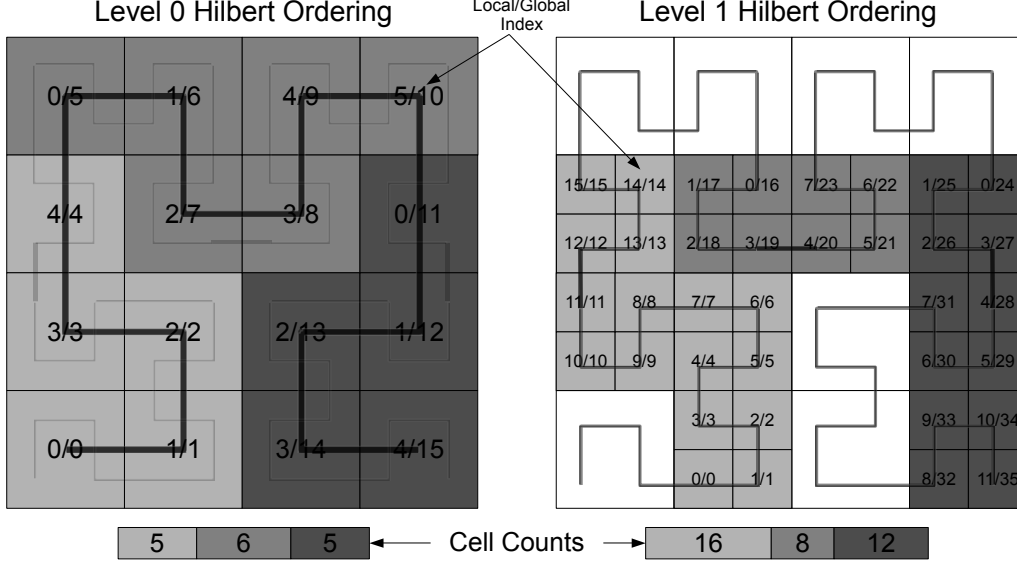
This approach permits a strict grouping of grids on processors by their sort order. This is true to at least the extent that grids obey the same type of child order inheritance true for individual cells. That is, while the order of two child cells will always be the same as that of their parent cells, it may not hold for grids. Cells have a unique distance along the space filling curve, however grids contain many cells each with different distances. A distance for the grid along the curve can be approximated by averaging that of all of its cells, or some subset of cells that appropriately samples the grid (ie the four corners or the center most subset). However this fuzziness in grid distances can result in child grids occasionally being ordered differently then their respective parents. However, this does not appear to be a major problem and would only result in slight increases in neighbor communication.

5.2. Hilbert Splitting

Occasionally splitting of a child grid into various pieces to accommodate its processor allocations $\mathcal{D}_{p'}^i$ is desired and there are two such instances in AstroBEAR. First, since most of the workload resides on the finest level grids, imbalances can sometimes be too large to be compensated for by the coarser threads. For this reason AstroBEAR always splits the finest level grids when necessary to achieve global load balancing. In general the number of finest level grids split will be of the order of the number of processors. Since the splitting occurs on the finest level, such fragmentation will not result in subsequent artificial fragmentation of higher level grids. Second, since the base grid in AMR simulations can be quite large it can often contain more then a single processor's share of the entire workload. For this reason the base grid is also split among the processors. The algorithm for splitting begins by taking a single grid and its non-zero processor allocations $\mathcal{D}_{p'}^i > 0$ that were determined by locally partitioning the set of $\delta_{p'}^p$ over the individual child workloads \mathcal{C}_l^i . This then gives the share of each individual child grid that should be assigned to each child processor. It then attempts to split the grid into pieces each with a workload equal to the processor allocations $\mathcal{D}_{p'}^i$. Additionally it attempts to construct the pieces so that they are properly ordered along the Hilbert curve. This is done through a recursive bisection algorithm described below:

1. Divide the list of weights into two pieces as equal as possible. For example if the list of weights were [.1 .2 .1 .3 .3] they would be split into [.1 .2 .1] and [.3 .3]
2. Determine the two possible split points along each dimension that break the grid into two proportionate pieces (ie either [.4 .6] or [.6 .4]).

Figure 7: Plot showing Hilbert ordering of level 0 and level 1 cells. Three colors correspond to processor containing the level 0 cells and the parent of the new level 1 cells.



3. For each of the possibilities evaluate the combined advance costs of the resulting grids. This will in general favor splits along the longest grid dimension.
4. If several possible splits have comparable combined advance costs then select the one that results in the largest difference in space-filling values between the two pieces with the correct sign.
5. If there were only two weights then we are done. Otherwise, continue to recursively bisect the left and right pieces.

For a large fixed grid simulation, the above algorithm applied to a square base grid in 2D (or a cube in 3D) on 4^n processors (8^n in 3D), will yield a distribution that traces out a level n Hilbert curve.

6. Performance Results

It is challenging to characterize the performance of an AMR code in general terms. This is because the performance on any given problem will depend on many factors that are specific to that problem such as the number of levels of refinement, the number of cells on each level (filling fractions), their spatial distribution, and the number of processors used. Here we give a few examples to demonstrate the complexity of the problem.

First consider a fixed grid 16x16 problem run on 5 processors. Ideally the 16x16 grid could be split into 5 even rectangular pieces to distribute among 5 processors. However, the best arrangement still leaves one processor with 55 cells which is 7% more than the ideal average of 51.2 effectively limiting the performance to 93%. We could split the grid into more than 5 pieces so that the maximum number of cells on any processor is closer

to the mean, but the additional overhead with having smaller grids would likely defeat the purpose.

The same problem however can be easily broken into 16 even pieces (each 4×4) to be perfectly distributed among 16 processors. However, running the problem on 16 processors instead of 1 will not generally allow the problem to complete in $1/16$ of the time for two primary reasons. First, before each 4×4 grid can update, it must share ghost data with the surrounding grids and there is some network latency in transmitting ghost data between grids. Additionally, while updating the 4×4 grid, calculations will be done using the ghost cells that will be repeated on the neighboring processors. This redundant calculation is not required if there is only one 16×16 grid. Thus updating 16 4×4 grids requires more calculations than updating 1 16×16 grid, so the CPU time will in general increase as more and more processors are used for the same problem even if the geometry of the problem allows for even distribution.


Now consider the same 16×16 grid on 16 processors and let's assume that there are 16 refined regions (each 2×2) that just happen to lie in the center of each 4×4 grid. These child grids will not be adjacent to other child grids and will only need to communicate with their parent grid. If each child grid is distributed on the same processor containing the parent, then the communication pattern is the same as that of the fixed grid run and the performance should actually be better because each processor now has more work that can be done independently. If however, instead of having 16 2×2 regions marked for refinement, the central 8×8 region was marked for refinement, then the communication pattern of the level 1 grids would be very similar to that of the level 0 although there would now be some additional communication between parent and child grids.

Thus the simplest and fairest test of an AMR code is to keep the number of cells per processor fixed as well as keeping the refined regions spatially connected. Furthermore, if we keep the filling fractions at $1/4$ for 2D ($1/8$ for 3D) then each level will have the same number of cells and the number of cells per level per processor will be a constant. The first problem we chose to embody this test philosophy was the advection of a power law density distribution in pressure equilibrium with its surroundings. The advection carried the density feature diagonally around a periodic square grid. The density profile and the refinement tolerances were chosen to keep the filling fraction at 25% for every level so that the number of cells per processor per level would be constant. We then ran a suite of simulations in which we varied the number of processors, the number of levels of AMR, and the number of cells per processor per level using both the serial AMR approach and the threaded algorithm.

Figure 8 shows the comparison between the serial and threaded AMR algorithms with maximum refinement levels of 0, 1, 2, & 4. Notice that both the serial and threaded versions give the same performance for a fixed grid (max-level = 0) calculation. This is because with only one level of AMR, the serial and threaded algorithms are the same. For max levels of 1, 2, & 4, the serial algorithm does not scale as well as the threaded, although the performance is independent of the number of levels of AMR. The threaded algorithm, however performs better with more levels of AMR. This is because there now exist multiple threads of execution which keep the processors from being idle.

In figure 9 we show the performance results for the serial and threaded algorithms keeping the number of levels of refinement fixed at 4, but varying the number of cells per processor per level from 8^2 to 256^2 . Strong scaling results are also shown (for a base grid of 256^2) by keeping the number of cells per processor level inversely proportional to

Figure 8: Plot showing performance comparison between serial and threaded AMR with maximum refinement levels of 0,1,2, & 4 and with 128^2 cells per processor per level.



the number of processors. Again we see that the threaded implementation out performs the serial implementation by 20% - 50%.

Figure 10 shows the ratio of memory used to store the distributed tree to the memory that would be required to store the global tree on each processor. Also shown are the ratio of total memory required to store the tree and data for the distributed tree algorithm compared to storing the global tree. As the grid sizes become larger, the ratio of the tree data to the grid data becomes smaller and the distributed tree becomes less important (at least until you reach larger numbers of processors).

We also tested the performance of the code out to 10^3 processors in a problem that engaged more of AstroBEAR's multi-physics capacities by advecting a magnetized cylinder across the domain until it was displaced by 1 cylinder radius. In this problem the size of the cylinder was chosen to give a filling fraction of approximately 12.5%. Thus in these AMR runs, the workload for the first refined level was comparable to the base level. The resolution of the base grid was adjusted to maintain 64^3 cells per processor. Figure 11 demonstrates that weak scaling for both fixed grid and for AMR versions of AstroBEAR 2.0 is reasonable out to 2048 processors. These results are consistent with what we found for both the memory and threading tests discussed above and indicate that the algorithm presented in this paper allows for efficient parallelization of multi-physics AMR simulations out to large numbers of processors.

Figure 9: Plot showing performance comparison between serial and threaded AMR where the number of cells per processor per level ranges from 8^2 to 256^2 . Note the scaling improves as larger and larger grids are used.




Figure 10: Plot showing memory savings of the distributed tree algorithm for the tree memory as well as for the memory required for tree and data.





Figure 11: Weak scaling for fixed grid and AMR versions of AstroBEAR out to 2048 processors.



7. Conclusion

As discussed in the introduction Adaptive Mesh Refinement methods were designed to allow high resolution simulations to be carried out at low computational cost. Highly parallel systems and their algorithms carried the same promise. The parallelization of AMR algorithms is, however, not straight forward however and there have been a number of different approaches [6, 9, 8, 1, 2], to solve the problem. While parallelization of a uniform mesh demands little communication between processors, AMR methods can demand considerable communication to maintain data consistency across the unstructured mesh as well as shuffling new grids from one processor to another to balance workload.

In this paper we have described our attempt to design and implement a new strategy for AMR parallelization with an eye towards running codes on large machines with $> 10^3$ processors. We have found that a threaded approach to the AMR algorithm significantly improves performance by allowing processors to remain busy while waiting for messages as well as to dynamically adjust distributions based on the progress of ongoing coarser grid updates. We have also shown that a distributed tree algorithm significantly reduces the amount of memory required (and corresponding communication) for simulations run on many processors with modest grid sizes.

We believe this threaded, distributed tree approach described in this paper holds considerable promise as a methodology for implementing AMR on ever-larger numbers of processors. Future work will entail elaborating refinements to the methods as we experiment with its implementation on different classes of problems and machines.

References

- [1] Chombo library.

Child iteration	Node's child's neighbors	Node's child's preceding overlaps	Node's child's succeeding overlaps
1	Node's neighbors' 1 st children	Node's preceding overlaps' R th children	Node's neighbor's 2 nd children
2	Node's neighbors' 2 nd children	Node's neighbors' 1 st children	Node's neighbors' 3 rd children
3	Node's neighbors' 3 rd children	Node's neighbors' 2 nd children	Node's neighbors' 4 th children
...
R-1	Node's neighbors' R th -1 children	Node's neighbors' R th -2 children	Node's neighbors' R th children
R	Node's neighbors' R th children	Node's neighbors' R th -1 children	Node's succeeding overlaps' 1 st children

- [2] A.M and Khokhlov. Fully threaded tree algorithms for adaptive refinement fluid dynamics simulations. *Journal of Computational Physics*, 143(2):519 – 543, 1998.
- [3] Marsha J Berger and Joseph Oliger. Adaptive mesh refinement for hyperbolic partial differential equations. *Journal of Computational Physics*, 53(3):484 – 512, 1984.
- [4] M.J. Berger and P. Colella. Local adaptive mesh refinement for shock hydrodynamics. *Journal of Computational Physics*, 82(1):64 – 84, 1989.
- [5] A. J. Cunningham, A. Frank, P. Varnière, S. Mitran, and T. W. Jones. Simulating Magnetohydrodynamical Flow with Constrained Transport and Adaptive Mesh Refinement: Algorithms and Tests of the AstroBEAR Code. *ApJS*, 182:519–542, June 2009.
- [6] P. MacNeice, K. M. Olson, C. Mobarry, R. de Fainchtein, and C. Packer. PARAMESH: A parallel adaptive mesh refinement community toolkit. *Computer Physics Communications*, 126:330–354, April 2000.
- [7] J. Mellor-Crummey, D. Whalley, and K. Kennedy. Improving memory hierarchy performance for irregular applications. *International Journal of Parallel Programming*, 29(3), June 2001.
- [8] B. W. O'Shea, G. Bryan, J. Bordner, M. L. Norman, T. Abel, R. Harkness, and A. Kritsuk. Introducing Enzo, an AMR Cosmology Application. *ArXiv Astrophysics e-prints*, March 2004.
- [9] U. Ziegler. The NIRVANA code: Parallel computational MHD with adaptive mesh refinement. *Computer Physics Communications*, 179:227–244, August 2008.

Appendix .1. Static Distributions for a Static Mesh

Consider what happens if the mesh is static so that the average workload per level $\overline{g_l}$ remains constant. Initially there will be a desired set of distributions d_l^p that will result in an actual set of distributions g_l^p . On the finest level they will agree, but on coarser levels there will be some difference between d_l^p and g_l^p . It is desirable that each successive desired distribution should match the original so that each successive actual distributions will also remain constant and grids will not be unnecessarily shuffled around. For example, consider the desired distribution of the highest level L after the first level L advance. Assuming the highest level workload follows the desired workload, $g_L^p = d_L^p = \overline{g_L} - \frac{\eta_L^p - \overline{\eta_L}}{s_L}$. The amount of time each processor will have waiting for every other processor to complete the level L advance is just $\max_p(g_L^p) - g_L^p$. If this time is spent advancing coarser grids,

$$\text{then the new sum of work done on coarser grids } \sum_{l'=0}^{L-1} w_l^{p*} = \sum_{l'=0}^{L-1} w_l^p + \left[\max_p(g_L^p) - g_L^p \right].$$

This gives a new predicted remaining work $\eta_L^{p*} = \eta_L^p - \left[\max_p(g_L^p) - g_L^p \right]$ and a new desired distribution

First Step		
1	<i>Receive</i> new grids & nodes along with their parents, neighbors, and preceding overlaps from parent processors	<i>Receive</i> succeeding overlaps from parent processors
2	Create new children and determine on which child processors they will go	
3	Determine which remote preceding nodes might have children that would overlap with its own children and <i>send</i> the relevant children info	Determine which remote succeeding nodes might have created children that would overlap with its own children and <i>send</i> the relevant children info
4	Determine which remote neighboring nodes might have children that would neighbor its own children and <i>send</i> the relevant children info	
5	Determine which local neighboring nodes have neighboring children	
6	Determine which local preceding nodes have children that overlap with its own	Determine which local succeeding nodes have children that overlap with its own
7	<i>Receive</i> new children from remote neighboring nodes and determine which of the neighbors' children neighbors its own children	
8	<i>Receive</i> children from remote preceding nodes and determine which of the nodes children overlaps with its own	<i>Receive</i> children from remote succeeding nodes and determine which of the nodes children overlaps with its own.
9	For each remote child, <i>send</i> the child grid's data as well as information about its parents, neighbors, & preceding overlaps.	For each remote child, <i>send</i> the child's succeeding overlaps.
Successive Steps		
10	Create new children and determine on which child processor they will go	
11	Determine which remote neighboring nodes might have old/new children that would overlap/neighbor its own new children and <i>send</i> the relevant children info	
12	Determine which local neighboring nodes might have old/new children that would overlap/neighbor its own new children	
13	<i>Receive</i> new children from remote neighboring nodes and determine which of the neighbors' children neighbors/overlaps its new/old children	
14	For each new remote child, <i>send</i> the child's information, and the information about its parent, neighbors, & preceding overlaps.	
15	For each old remote child, <i>send</i> the child's succeeding overlaps.	

Table .1: The split rows denote actions taken by the current iteration of nodes (left) and the old iteration of nodes (right). Otherwise the actions are taken only by the current iteration of nodes.

$$\begin{aligned}
d_L^{p*} &= \overline{g_L} - \frac{\eta_L^{p*} - \overline{\eta_L^{p*}}}{s_L^*} = \overline{g_L} - \frac{\eta_L^p - \left(\max_p(g_L^p) - g_L^p\right) - \left[\overline{\eta_L} - \left(\max_p(g_L^p) - \overline{g_L}\right)\right]}{s_L - 1} \\
&= \overline{g_L} - \frac{\eta_L^p - \overline{\eta_L} + g_L^p - \overline{g_L}}{s_L - 1} = \overline{g_L} - \frac{\eta_L^p - \overline{\eta_L}}{s_L} = d_L^p
\end{aligned} \tag{.1}$$

More generally we can consider the distribution of any given level l after completing c_l steps leaving $s_l = 2^l - c_l$ remaining steps.

$$d_l^p = \overline{g_l} - \frac{\eta_l^p - \overline{\eta_l}}{s_l}$$

If $\eta_l^p - \overline{\eta_l} \propto s_l$, then the desired distribution will remain fixed.

$$\eta_l^p - \overline{\eta_l} = \left(\sum_{l'=0}^{l-1} s_{l'} g_{l'}^p - w_{l'}^p \right) - \left(\sum_{l'=0}^{l-1} s_{l'} \overline{g_{l'}} - \overline{w_{l'}} \right) \tag{.2}$$

The workload completed on the coarser grids will be the difference between the maximum work completed on all of the previous grids by any processor and the local work completed on all of the previous grids. This then gives

$$\sum_{l'=0}^{l-1} w_{l'}^p = \max_p \left[\sum_{l'=0}^L (2^{l'} - s_{l'}) g_{l'}^p \right] - \sum_{l'=0}^L (2^{l'} - s_{l'}) g_{l'}^p \tag{.3}$$

and we have

$$\begin{aligned}
\eta_l^p - \overline{\eta_l} &= \left(\sum_{l'=0}^{l-1} s_{l'} g_{l'}^p + \sum_{l'=0}^L (2^{l'} - s_{l'}) g_{l'}^p \right) - \left(\sum_{l'=0}^{l-1} s_{l'} \overline{g_{l'}} + \sum_{l'=0}^L (2^{l'} - s_{l'}) \overline{g_{l'}} \right) \\
&= \sum_{l'=0}^L 2^{l'} (g_{l'}^p - \overline{g_{l'}}) + \sum_{l'=l}^L s_{l'} (\overline{g_{l'}} - g_{l'}^p)
\end{aligned} \tag{.4}$$

Now $\sum_{l'=0}^L 2^{l'} g_{l'}^p = \eta_{L+1}^p$ it the initial entire predicted workload and is balanced by the highest level distribution so the value on any processor equals the average and the first term cancels. The second term can be re-written by recognizing that at the moment we are distributing level l , for $l' \geq l$, $s_{l'} = 2^{l'-l} s_l$. This gives

$$\eta_l^p - \overline{\eta_l} = s_l \sum_{l'=l}^L 2^{l'-l} (\overline{g_{l'}} - g_{l'}^p) \propto s_l \tag{.5}$$

Therefore if the mesh is static, the distribution will remain static as well.